

The Foundations of Computational Imaging

A signal processing perspective



©SHUTTERSTOCK.COM/TRIFF

Twenty-five years ago, the field of computational imaging arguably did not exist, at least not as a standalone arena of research activity and technical development. Of course, the idea of using computation to form images had been around for several decades, largely thanks to the development of medical imaging—such as magnetic resonance imaging (MRI) and X-ray tomography—in the 1970s and synthetic-aperture radar (SAR) even earlier. Yet, a quarter of a century ago, such technologies would have been considered to be a subfocus of the wider field of image processing. This view started to change, however, in the late 1990s with a series of innovations that established computational imaging as a scientific and technical pursuit in its own right.

Introduction

In this article, we provide a signal processing perspective on the area of computational imaging, focusing on its emergence and evolution within the signal processing community. First, in the “Historical Perspective” section, we present an overview of the technical development of computational imaging wherein we trace the historical development of the field from its origins through to the present day. Then, in the “Signal Processing Society Involvement” section, we provide a brief overview of the involvement of the IEEE Signal Processing Society (SPS) in the field. Finally, in the “Conclusions” section, we make several concluding remarks.

Historical perspective

We begin our discourse by tracing the history of computational image formation. We start this historical perspective with its origins in physics-dependent imaging and then proceed to model-based imaging, including the impact of priors and sparsity. We next progress to recent data-driven and learning-based image formations, finally coming full circle back to how physics-based models are being merged with big data and machine learning for improved outcomes.

Computational imaging can be defined in a number of ways, but for the purposes of the present article, we define

it as the creation of an image from measurements wherein computation plays an integral role in the image-formation process. In contrast to historical “standard” imaging in which optics play the central role in image formation, in computational imaging, it is computation, in the form of algorithms running on computers, that assumes the primary burden of producing images. Well-known examples would include X-ray-based tomography, SAR, and MRI. In such cases, the data produced by the sensing instrument are generally not images, and thus, require processing to produce the desired useful output in the form of an image. To fix ideas and notation, we denote the unknown but desired image by $\mathbf{x} \in \mathbb{R}^N$, such that \mathbf{x} contains N pixels. In our problems of interest, we cannot directly observe \mathbf{x} , but instead observe a set of data $\mathbf{y} \in \mathbb{R}^M$, which has been measured through a process connected to a sensor C . This relationship can be represented mathematically by

$$\mathbf{y} = C(\mathbf{x}) + \mathbf{n} \quad (1)$$

where \mathbf{n} is a noise signal. The goal of computational imaging is then to estimate the image \mathbf{x} from knowledge of both the data \mathbf{y} as well as the imaging system or measurement process, C , i.e., it naturally involves solving an inverse problem. Several examples of computational imaging are depicted in Figure 1, illustrating their sensing process and resulting images.

For many common imaging problems, C in (1) is (or can be well approximated by) a linear operator, i.e., a matrix, such that

$$C(\mathbf{x}) = C\mathbf{x}, \quad C \in \mathbb{R}^{M \times N}. \quad (2)$$

In this case, when C is invertible and well conditioned, the inverse problem is mathematically straightforward, although even in this case, it is often computationally challenging due to the size of the problem. For example, an image of modest size—say, $1,024 \times 1,024$ —corresponds to an \mathbf{x} with 1 million variables and a C with 10^{12} elements. The inverse problem becomes even more difficult when C is an ill-conditioned matrix or a nonlinear operator or produces an observation \mathbf{y} with fewer samples than the number of unknowns in \mathbf{x} .

In the sequel, we provide a roughly chronological road map of computational imaging from a signal processing perspective. Four broad domains, as illustrated in Figure 2, will be discussed in turn: physics-driven imaging, model-based image reconstruction (MBIR), data-driven models and learning, and learning-based filters and algorithms. Our discussion concludes with current-day algorithmic methods that effectively join the right side of Figure 2 back up to the left, i.e., techniques that couple physical models with learned prior information through algorithms.

Physics-driven imaging: Explicit inversion

We first focus on physics-driven imaging, wherein images are created based on idealized imaging models. The idea here is to form images using the physical inversion operator with the key challenge being the design of efficient algorithms for the calculation of $\hat{\mathbf{x}} = C^{-1}(\mathbf{y})$ as an approximation to the desired

image \mathbf{x} in (1). When data are plentiful and of high quality, and the system is designed to closely approximate certain simple classes of sensing operators, this direct-inversion approach can work reasonably well. Such image-formation methods often represent the first approach taken in the development of a new imaging modality and will often invoke very little (or very rudimentary) prior information, if at all. We discuss four example modalities in greater detail next.

Analog cameras

In a film camera, lenses are used to bend the light and focus it onto a film substrate, as depicted in Figure 3(a). Digital cameras work largely the same way by simply placing a digitizer at the film plane. In terms of (1)–(2) then, this example has an ideal linear model with $C = I$ so that the final image is really just the observation, assuming negligible noise. Traditionally, the primary path to improving the image quality was through improvements to the optical path itself, that is, through the use of better lenses that bring the physical sensing process closer to the ideal identity $C = I$.

X-ray tomography

X-ray tomography is used in applications such as nondestructive evaluation, security, and medical imaging and is, in essence, an image of the attenuation produced by an object as X-rays pass through it, as illustrated in Figure 3(b). The negative log ratio of the attenuated output to the input incident energy is taken as the observation, with a simplified physical model being

$$\mathbf{y}(L) = \int_L \mathbf{x}(s) ds \quad (3)$$

where L is a given ray path. Mathematically, the observation (or projection), $\mathbf{y}(L)$, is a line integral of the desired attenuation image $\mathbf{x}(s)$ along the line L . The collection of all such projections for all lines L (i.e., in every direction) defines what is called the *Radon transform* of the image \mathbf{x} . The Radon transform is a linear operation and, assuming no noise, can be represented by (1), with C in this case being defined by the integral operator in (3). An explicit analytic inverse of the Radon transform (i.e., C^{-1}) exists and forms the basis of the image-formation method used by commercial scanners. This inversion approach, called the *filtered back-projection (FBP) algorithm*, is very efficient in both computation and storage, requiring only simple filtering operations on the projections followed by an accumulation process wherein these filtered projections are summed back along their projection directions [4]; however, it assumes the existence of an infinite

In this article, we provide a signal processing perspective on the area of computational imaging, focusing on its emergence and evolution within the signal processing community.

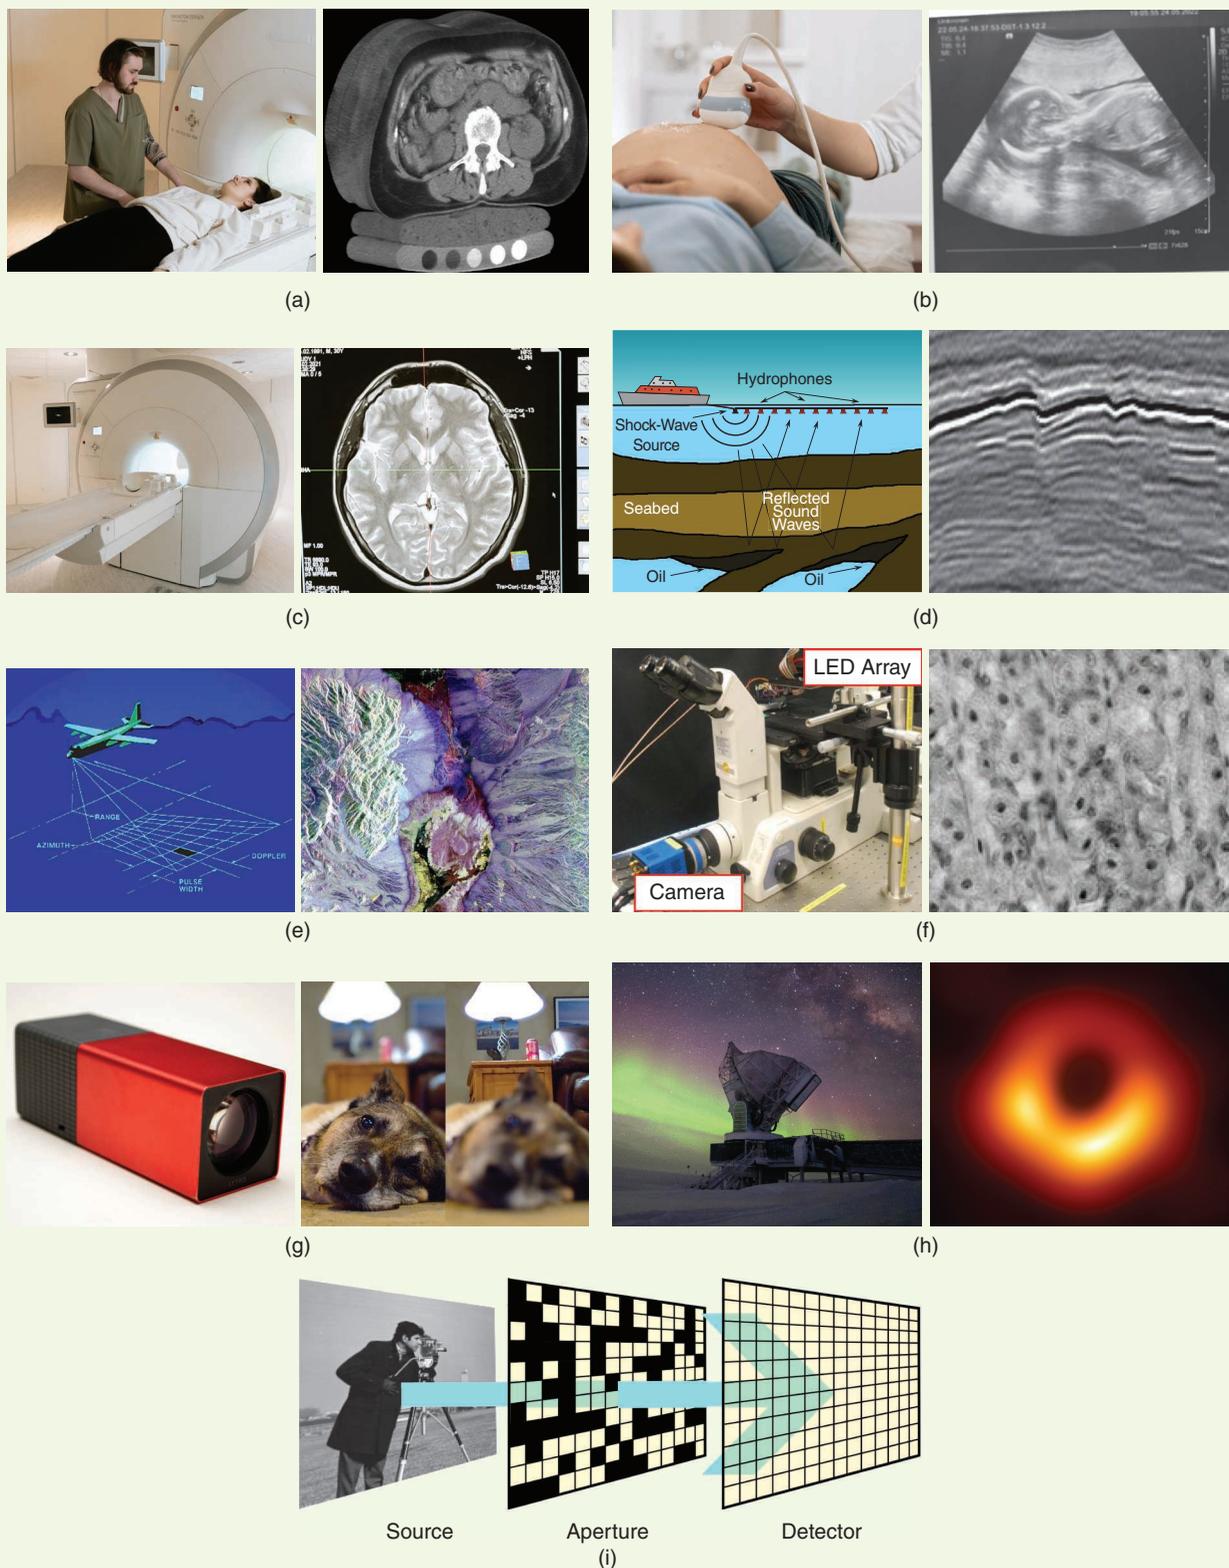


FIGURE 1. Examples of computational-imaging modalities. (a) X-ray tomography. (b) Ultrasound imaging. (c) MRI. (d) Seismic imaging. (e) Radar imaging. (f) Computational microscopy. (g) Light-field imaging. (h) Astronomical imaging. (i) Coded-aperture imaging. Sources: (a) Mart Production (left); MindwaysCT Software, CC BY-SA 3.0 (right). (b) Mart Production (left); Daniel Kondrashin (right). (c) Mart Production. (d) Adapted from [1] (right). (e) Wclxs11, CC BY-SA 3.0 (left); NASA/JPL (right). (f) Adapted with permission from [2]; ©The Optical Society. (g) Dcoetzee, CC0 (left); Doodybutch, CC BY-SA 4.0 (right). (h) Adapted from [3].

continuum of projections, which is not possible in practice. Nevertheless, excellent reconstructed images are possible if the sampling of $\mathbf{y}(L)$ is sufficiently dense. Thus, higher quality images are obtained by making the X-ray machine better approximate the underlying assumptions of the Radon-transform inversion.

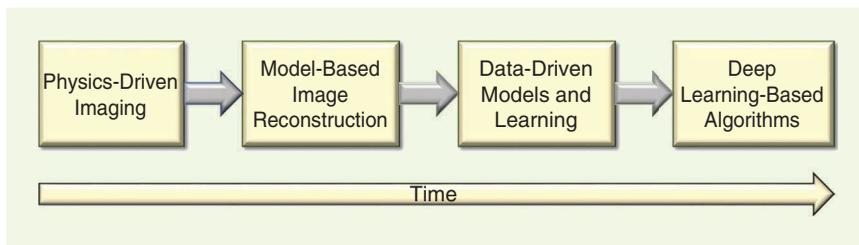


FIGURE 2. An overview of the historical evolution of computational imaging.

MRI

As depicted in Figure 3(c), MRI images the anatomy of the body through radio-frequency excitation and strong magnetic fields—unlike X-ray imaging, no ionizing radiation is used. Classical MRI acquisition is usually modeled as producing observations according to the equation

$$\mathbf{y}(f) = \int \mathbf{x}(s) e^{-i2\pi s f} ds \quad (4)$$

where it can be seen that the observations $\mathbf{y}(f)$ are values of the desired image in the Fourier domain. The basic MRI acquisition acquires samples of these Fourier values line by line in the Fourier space, called the k -space, and once sufficient Fourier samples are obtained, an image is produced by the application of an inverse Fourier transform. As with X-ray tomography, the image formation follows from an analytic formula for the inverse of the Fourier-based sensing operator such that improved imagery is obtained through the denser and more complete sampling of the k -space [5].

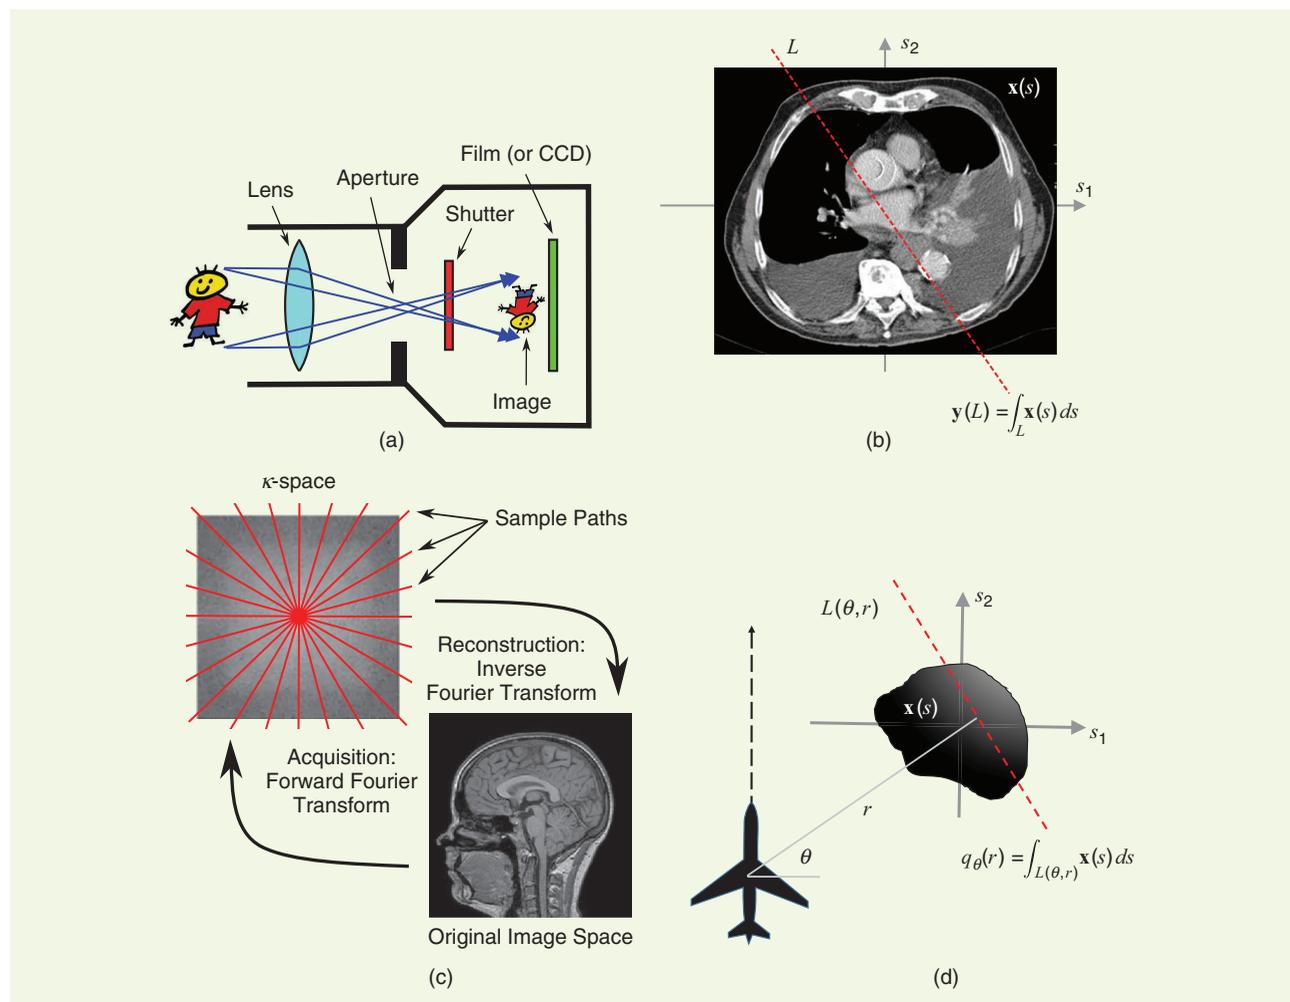


FIGURE 3. Four common forms of physics-driven imaging with explicit inversions. (a) Analog camera. (b) X-ray tomography. (c) MRI. (d) SAR. CCD: charge-coupled device. Sources: (b) Image from Nevit Dilmen, CC BY-SA 3.0. (c) MRI image from the IXI Dataset (<https://brain-development.org/ixi-dataset/>), CC BY-SA 3.0.

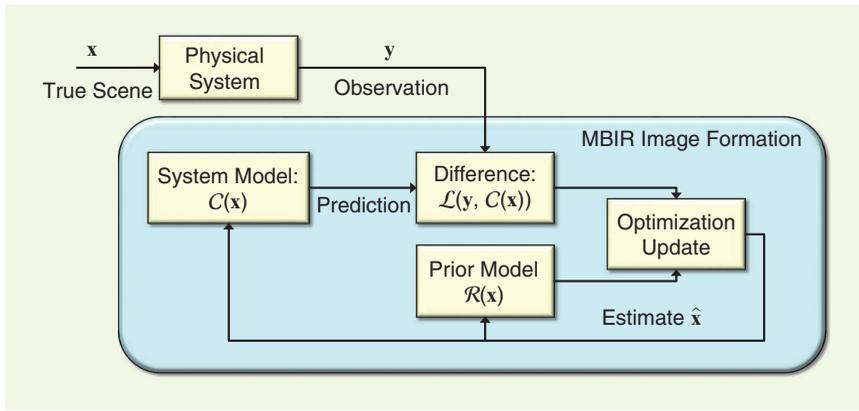


FIGURE 4. A general framework for MBIR.

Spotlight-mode SAR

Spotlight-mode SAR, as shown in Figure 3(d), is able to create high-resolution images of an area day or night and in all types of weather and is thus widely used in remote sensing applications. SAR works by transmitting microwave chirp pulses toward the ground and has a resolution that is largely independent of the distance to the region of interest. SAR finds use in mapping and monitoring vegetation and sea ice and in NASA planetary missions as well as in military applications, just to name a few. The SAR data-acquisition process can be modeled as

$$y_{\theta}(t) = \int q_{\theta}(r) e^{-j\Omega(t)r} dr \quad (5)$$

where $\Omega(t)$ is a time-dependent “chirped” frequency, and

$$q_{\theta}(r) = \int_{L(\theta,r)} \mathbf{x}(s) ds \quad (6)$$

is the projection of the scattering field $\mathbf{x}(s)$ along the line L at angle θ and range r such that (5) is a 1D Fourier transform of the projection along the range direction. Thus, the observations for SAR are again related to the values of the desired image $\mathbf{x}(s)$ in the Fourier domain, similar to MRI. Combining (5) and (6), one can show that these observations are Fourier values of $\mathbf{x}(s)$ on a polar grid; consequently, the standard image-formation algorithm for such SAR data is the polar-format algorithm [6], which resamples the acquired polar Fourier data onto a rectangular grid and then performs a standard inverse Fourier transform on the regridded data. As in our other examples, the image-formation process follows from an analytic formula for the inverse of the (Fourier-based) sensing operator, and improved imagery is obtained by extending the support region of the acquired Fourier samples, which are related to both the angular sector of SAR observation (related to the flight path of the sensor) as well as the bandwidth of the transmitted microwave chirp.

In all the examples that we have just discussed, image formation comprises analytic inversion following from the physics of the measurement operator $[C$ in (1)]. These inversion approaches

operate on the observed data, but the algorithms themselves are not dependent on the data; i.e., the structure of the algorithm, along with any parameters, is fixed at the time of the algorithm design based exclusively on the inversion of the measurement operator and is not learned from image data. These examples of the early period of computational imaging can be characterized in our taxonomy by (comparatively) low computation (e.g., Fourier transforms) and the presence of “small-data” algorithms (i.e., just the observations). When data are complete and of high quality and the system is designed to closely approximate the

assumptions underlying the inversion operators, these approaches can work very well. Yet, the physical systems are constrained by the algorithmic assumptions on which the image formation is built, and if the quality or quantity of observed data is reduced, much of the corresponding imagery can exhibit confounding artifacts. An example would be standard medical tomography—since the system will create an image using the FBP, the system needs to be such that a full 180° of projections are obtained at a sufficient sampling density.

MBIR: The rise of computational imaging

Computational imaging really flourished during the next phase we consider. This phase has been called *MBIR* and, in contrast to the situation discussed in the “Physics-Driven Imaging: Explicit Inversion” section, is characterized by the use of explicit models of both sensing physics as well as image features. The major conceptual shift was that image formation was cast as the solution to an optimization problem rather than as the output of a physically derived inversion algorithm. Next, we present the specifics of this optimization problem, consider the role of prior models in its formulation, and explore several algorithms for its solution.

Image formation as optimization with explicit models

In general, in MBIR, the image-formation optimization problem can be taken to be of the form

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{L}(\mathbf{y}, C(\mathbf{x})) + \mathcal{R}(\mathbf{x}) \quad (7)$$

where $\hat{\mathbf{x}}$ is the resulting output image (an approximation to the true image); \mathcal{L} is a loss function that penalizes any discrepancy between the observed measurement \mathbf{y} and its prediction $C(\mathbf{x})$; and \mathcal{R} is a regularization term that penalizes solutions that are unlikely according to prior knowledge of the solution space. This optimization is depicted schematically in Figure 4. Arguably, the simplest example of this formulation is Tikhonov regularization

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{y} - C(\mathbf{x})\|_2^2 + \lambda \|\Gamma(\mathbf{x})\|_2^2 \quad (8)$$

where λ is a parameter controlling the amount of regularization, and Γ is a problem-specific operator [7], often chosen to be a derivative. This formulation also connects to probabilistic models of the situation. In particular, the first term in (8) can be associated with a log-likelihood $\log p(\mathbf{y}|\mathbf{x}, C)$ and the second term with a log prior $\log p(\mathbf{x})$ under Gaussian assumptions. With this association, (8) represents a maximum a posteriori (MAP) estimate of the image given a likelihood and a prior, i.e.,

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmax}} p(\mathbf{x}|\mathbf{y}, C) = \underset{\mathbf{x}}{\operatorname{argmax}} \log p(\mathbf{y}|\mathbf{x}, C) + \log p(\mathbf{x}). \quad (9)$$

There are a number of major advantages gained from the conceptual shift represented by viewing image formation as the solution of (7). One advantage is that this view separates out the components of the formulation from the algorithm used to solve it; i.e., the overall problem is partitioned into independent modeling and optimization tasks. Indeed, there are many approaches that can be used to solve (7), allowing discussion of the algorithm to be decoupled from the debate about the problem formulation [although obviously some function choices in (7) correspond to easier problems, and thus, simpler algorithms].

Another advantage of this explicit focus on models is that one can consider a much richer set of sensing operators since we are no longer limited to operators possessing simple, explicit, and closed-form inverse formulations. For example, in X-ray tomography, the FBP algorithm is an appropriate inversion operator only when a complete (uniform and densely sampled) set of high-quality projection data is obtained—it is not an appropriate approach if data are obtained, for example, over only a limited set of angles. On the other hand, the formulation (7) is agnostic to such issues, requiring only that the C operator accurately captures the actual physics of acquisition. Thus, model-based inversion approaches have been successfully applied in situations involving novel, nonstandard, and challenging imaging configurations that could not have been considered previously. Furthermore, one can now consider the joint design of sensing systems along with inversion, as occurs in computational photography.

A third advantage of model-based image formation is that (7) can be used to explicitly account for noise or uncertainty in the data—the connection to statistical methods and MAP estimation as alluded to previously [i.e., (9)] makes this connection obvious. For example, rather than using a loss function corresponding to a quadratic penalty arising from Gaussian statistics (as is common), one can consider instead a log-likelihood associated with Poisson-counting statistics, which arises naturally in photon-based imaging. The use of such models can provide superior noise reduction in low-signal situations.

A final advantage of model-based inversion methods is the explicit use of prior-image-behavior information, as

captured by the term $\mathcal{R}(\mathbf{x})$. We focus on the impact of prior models next.

Focus on prior models and the emergence of sparsity

The growth of model-based methods led to a rich exploration of choices for the prior term $\mathcal{R}(\mathbf{x})$ in (7). The simplest choice is perhaps a quadratic function of the unknown \mathbf{x} , as illustrated in (8). Such quadratic functions can be viewed as corresponding to a Gaussian assumption on the statistics of $\Gamma(\mathbf{x})$. While simple and frequently leading to efficient solutions of the corresponding optimization problem, the resulting image estimates can suffer from blurring and the loss of image structure as these types of priors correspond to aggressive global penalties on image or edge energy.

Such limitations led to a surge in the development and use in the late 1980s and 1990s of nonquadratic functions that share the property that they penalize large values less than the quadratic penalty does. Additionally, when applied to image derivatives, they promote edge formation. In Table 1,

we present a number of the nonquadratic penalty functions that arose during this period, separated into those that are convex and those that are not. (The penalties tabulated in Table 1 are functions on scalars. To form $\mathcal{R}(\mathbf{x})$, these scalar penalties could be applied element by element to

$\mathbf{x} = [x_1 \ x_2 \ \dots \ x_N]^T$, e.g., $\mathcal{R}(\mathbf{x}) = \sum_i \varphi(x_i)$. Alternatively, referring to (8), they could likewise be applied to elements of $\Gamma(\mathbf{x})$.) In general, convex functions result in easier optimization problems, while nonconvex functions possess more aggressive feature preservation at the expense of more challenging solution computation.

A key property that these nonquadratic penalties promoted was the sparsity of the corresponding quantity. In particular, when applied to \mathbf{x} itself, the resulting solution becomes sparse, or when applied to $\Gamma(\mathbf{x})$, the quantity $\Gamma(\mathbf{x})$ becomes sparse. A common example is to cast Γ as an approximation to the gradient operator such that the edge field then becomes sparse, resulting in piecewise constant (mainly flat) solutions [16]. Eventually, interest nucleated around this concept of sparsity and, in particular, on the use of the ℓ_0 and ℓ_1 norms as choices in defining $\mathcal{R}(\mathbf{x})$ (the ℓ_1 norm corresponds to the last row in Table 1) [17].

Table 1. A selection of nonquadratic prior penalties $\varphi(t)$.

Convex Penalties	Nonconvex Penalties
$ t $ [8]	$\frac{t^2}{1+t^2}$ [9], [10]
$\log \cosh t$ [11]	$\min\{t^2, 1\}$ [12]
$\min\{ t ^2, 2 t - 1\}$ [10], [13]	$\log(1+t^2)$ [14]
	$\frac{ t }{1+ t }$ [15]
$ t ^p, p \geq 1$ [16]	$ t ^p, 0 < p < 1$

One of the more visible applications of sparsity in model-based reconstruction is compressed sensing.

One of the more visible applications of sparsity in model-based reconstruction is compressed sensing (CS) [18], [19]. In brief, under certain conditions, CS permits the recovery of signals from their linear projections into a much lower dimensional space. That is, we recover \mathbf{x} from $\mathbf{y} = C\mathbf{x}$, where \mathbf{x} has length N , \mathbf{y} has length M , and C is an $M \times N$ measurement matrix with the subsampling rate (or substrate) being $S = M/N$ with $M \ll N$. Because the number of unknowns is much larger than the number of observations, recovering every $\mathbf{x} \in \mathbb{R}^N$ from its corresponding $\mathbf{y} \in \mathbb{R}^M$ is impossible in general. The foundation of CS, however, is that, if \mathbf{x} is known to be sufficiently sparse in some domain, then exact recovery of \mathbf{x} is possible. Such sparsity can be with respect to some transform T such that, when the transform is applied to \mathbf{x} , only $K < M \ll N$ coefficients in the set of transform coefficients $\mathbf{X} \equiv T\mathbf{x}$ are nonzero. Relating this situation back to (8), we can formulate the CS recovery problem as

$$\hat{\mathbf{X}} = \underset{\mathbf{X}}{\operatorname{argmin}} \|\mathbf{y} - CT^{-1}\mathbf{X}\|_2^2 + \lambda \|\mathbf{X}\|_p \quad (10)$$

with the final reconstruction being $\hat{\mathbf{x}} = T^{-1}\hat{\mathbf{X}}$. Ideally, for a T -domain sparse solution, we set $p = 0$ in (10), invoking the ℓ_0 pseudonorm, which counts nonzero entries. Since this choice results in an NP-hard optimization, however, it is common to use $p = 1$, thereby applying a convex relaxation of the ℓ_0 problem. Interestingly, if the underlying solution is sufficiently sparse, it can be shown that the two formulations yield the same final result [18], [19], [20]. Additionally, for exact recovery, it is sufficient that the image transform T and the measurement matrix C be “mutually incoherent” in the sense that C cannot sparsely represent the columns of the transform matrix T . Accordingly, an extensive number of image transforms T have been explored; additionally, large-scale images can be reconstructed by applying the formulation in a block-by-block fashion (e.g., [21]). CS has garnered a great deal of interest in the computational-imaging community in particular due to the demonstration of devices (e.g., [22]) that conduct the compressed signal-sensing process $\mathbf{y} = C\mathbf{x}$ entirely within optics, thereby acquiring the signal and reducing its dimensionality simultaneously with little to no computation.

Optimization algorithms

One of the challenges with MBIR-based approaches is that the resulting optimization problems represented by (7) must be solved. Fortunately, solutions of such optimization problems have been well studied, and a variety of methods exist, including the projected gradient-descent algorithm, the Chambolle-Pock primal-dual algorithm, and the alternating direction method of multipliers (ADMM) algorithm as well as others. These methods are, in general, iterative algorithms composed of a sequence of steps that are repeated until a stopping condition is reached. Of particular interest is the

ADMM algorithm as well as other similar methods exploiting proximal operators [23], [24], [25]. Such methods split the original problem into a series of pieces by way of associated proximal operators. Specifically, the ADMM algorithm for solving (7) recasts (7) with an additional variable \mathbf{z} and an equality constraint

$$\hat{\mathbf{x}} = \underset{\mathbf{x}, \mathbf{z}}{\operatorname{argmin}} \mathcal{L}(\mathbf{y}, C(\mathbf{z})) + \mathcal{R}(\mathbf{x}), \quad \text{such that } \mathbf{x} = \mathbf{z} \quad (11)$$

which is solved via the iterations

$$\begin{aligned} \mathbf{x}^{(k+1)} &\leftarrow \underset{\mathbf{x}}{\operatorname{argmin}} \frac{\rho}{2} \|\mathbf{x} - \mathbf{z}^{(k)} + \mathbf{u}^{(k)}\|_2^2 + \mathcal{R}(\mathbf{x}) \\ &= \operatorname{prox}_{\mathcal{R}, \rho/2}(\mathbf{z}^{(k)} - \mathbf{u}^{(k)}) \end{aligned} \quad (12)$$

$$\mathbf{z}^{(k+1)} \leftarrow \underset{\mathbf{z}}{\operatorname{argmin}} \frac{1}{2} \mathcal{L}(\mathbf{y}, C(\mathbf{z})) + \frac{\rho}{2} \|\mathbf{x}^{(k+1)} - \mathbf{z} + \mathbf{u}^{(k)}\|_2^2 \quad (13)$$

$$\mathbf{u}^{(k+1)} \leftarrow \mathbf{u}^{(k)} + \mathbf{x}^{(k+1)} - \mathbf{z}^{(k+1)} \quad (14)$$

where \mathbf{u} is the scaled dual variable, and ρ is a penalty parameter. We have indicated previously that (12) is a proximal operator, effectively performing smoothing or denoising to its argument. We will return to this insight later as we consider

the incorporation of learned information into computational imaging. Note that the ADMM algorithm comprises an image-smoothing step (12), a data- or observation-integration step (13), and a simple reconciliation step (14).

In the model-based approach to image reconstruction discussed in this section,

image formation is accomplished through the solution of an optimization problem, and underlying models of acquisition and image are made explicit. The prior-image models can serve to stabilize situations with poor data, and conversely, the observed data can compensate for overly simplistic prior-image models. These MBIR methods, including the use of nonquadratic models and the development of CS, have had a profound impact on the computational-imaging field. They have allowed the coupled design of sensing systems and inversion methods wherein sensor design can be integrated with algorithm development in ways not previously possible. The impact has been felt in fields as disparate as SAR, computed tomography, MRI, microscopy, and astronomical science. These methods are characterized in our taxonomy by relatively high computation (resulting from the need to solve relatively large optimization problems iteratively) and the use of “small-data” algorithms (again, just the observations).

Data-driven models and learning: Dictionaries

The next phase we consider is that in which data start to take on an important role in modeling, with approaches characterized by the increasing use of models derived from data rather than physics or statistics. This increased focus on data in modeling rather than on analytic models provides the

Perhaps the simplest data-driven extension of the MBIR paradigm can be found in the development of dictionary learning.

opportunity for an increase in explanatory richness. Perhaps the simplest data-driven extension of the MBIR paradigm can be found in the development of dictionary learning [17], [26], [27], [28]. The idea behind dictionary learning is that a noisy version of a signal \mathbf{x} can be approximated by a sparse linear combination of a few elements of an overcomplete dictionary D . This problem can be cast as an MBIR-type formulation as

$$\hat{\alpha} = \underset{\alpha}{\operatorname{argmin}} \|\mathbf{x} - D\alpha\|_2^2 + \lambda \|\alpha\|_0 \quad (15)$$

wherein the dictionary D has, in general, many more columns than rows, and thus, enforcing sparsity selects the most important columns. The final estimate is obtained as $\hat{\mathbf{x}} = D\hat{\alpha}$.

Data-based learning is introduced into this framework by employing a large set of training samples $\{x_1, x_2, \dots, x_N\}$ to learn the dictionary D . This dictionary-learning process can be cast conceptually as another MBIR-style sparsely constrained optimization

$$\hat{D}, \{\hat{\alpha}_i\} = \underset{D, \{\alpha_i\}}{\operatorname{argmin}} \sum_{i=1}^N \|\mathbf{x}_i - D\alpha_i\|_2^2 + \lambda \|\alpha_i\|_0 \quad (16)$$

$$\| [D]_j \|_2 \leq 1$$

where $\| [D]_j \|_2$ denotes the norm of column j of D , this norm constraint being necessary to avoid unbounded solutions for D . Ultimately, this formulation uses training data as well as a sparsity constraint to learn an appropriate representational framework for a signal class, and the resulting dictionary can then be used as a model for such signals in subsequent reconstruction problems. We note that while (16) conveys the aim of dictionary learning, in practice, a variety of different formulations are used.

Learning-based filters and algorithms: The appearance of deep learning

Renewed interest in the application of neural networks (NNs) to challenging estimation and classification problems occurred

when AlexNet won the ImageNet Challenge in 2012. The availability of large datasets combined with deep convolutional NNs (CNNs) as well as advanced computing hardware such as graphics processing units (GPUs) has enabled a renaissance of NN-based methods with outstanding performance and has created a focus on data-driven methods. Beyond classification, deep CNNs have achieved state-of-the-art performance in computer-vision tasks ranging from image denoising to image superresolution. Naturally, these deep learning models have made their way into the world of computational imaging in a variety of ways, several of which we survey next.

Estimate post-processing

Perhaps the simplest way of folding deep learning into computational-imaging problems is to apply a deep learning-based image enhancement as a post-processing step after a reconstructed image has been formed. In doing so, one uses an existing inversion scheme—such as FBP for X-ray tomography or inverse Fourier transformation for MRI—to create an initial reconstructed image. A deep network is then trained to bring that initial estimate closer to a desired one by removing artifacts or noise. The enhancement can be done directly on the formed image, or more recently, the deep enhancement network is trained on a set of residual images between initial estimates and high-quality targets. Approaches in this vein are perhaps the most straightforward way to include deep learning in image formation and were thus some of the first methods developed. Example application domains include X-ray tomography [29], [30] with subsampled and low-dose data as well as MRI [29], [31] with subsampled Fourier data. Figure 5 illustrates this learning-driven post-processing approach for subsampled MRI along with a physics-driven explicit inversion as well as an MBIR-based reconstruction.

Data preprocessing

Another use of deep learning in computational imaging is as a preprocessing step. In these methods, learning is used to “correct” imperfect, noisy, or incomplete data. Accordingly,

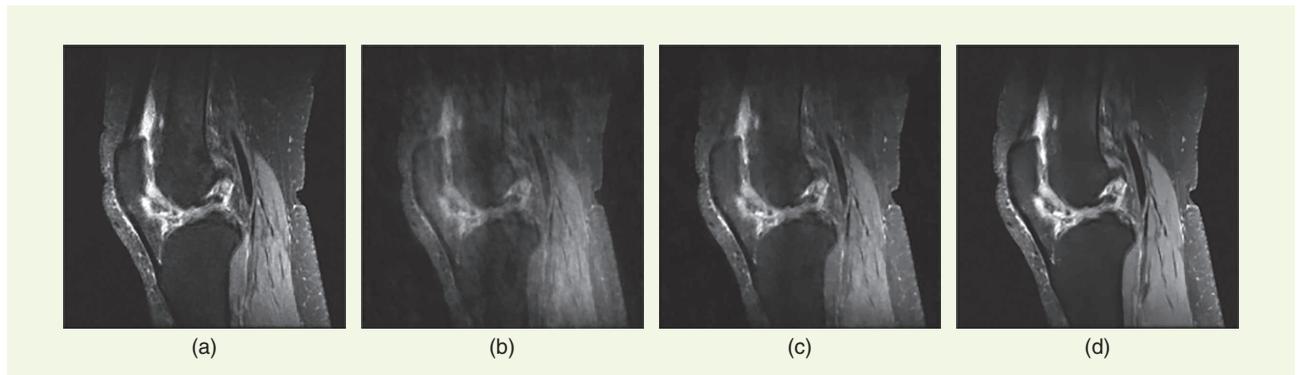


FIGURE 5. Subsampled MRI reconstruction and performance in terms of signal-to-noise ratio (SNR). (a) Inverse Fourier transform of complete data. (b) Reconstruction of sixfold subsampled data via the inverse Fourier transform as described in the “MRI” section; SNR = 11.82 dB. (c) Reconstruction of sixfold subsampled data via a convex optimization regularized with a total-variation criterion, i.e., an MBIR reconstruction in the form of (7); SNR = 15.05 dB. (d) Reconstruction via the inverse Fourier transform of the subsampled data followed by post-processing enhancement with a CNN as described in the “Estimate Post-processing” section; SNR = 17.06 dB. (Source: Adapted from [29], ©2017 IEEE.)

data can be made to more closely match the assumptions that underlie traditional physics-based image-formation methods, which can then be used more effectively. Consequently, such an approach can leverage existing conventional workflows and dedicated hardware. One example can be found in X-ray tomography [32] wherein projection samples in the sinogram domain that are missed due to the sparse angular sampling used to reduce dose are estimated via a deep network; afterward, these corrected data are used in a conventional FBP algorithm.

Learned explicit priors

Yet another use of deep learning in computational imaging has been to develop explicit data-derived learned priors that can be used in an MBIR framework. An example of this approach can be found in [33] wherein a K -sparse patch-based autoencoder is learned from training images and then used as a prior in an MBIR reconstruction method. An autoencoder (depicted in Figure 6) is a type of NN used to learn efficient reduced-dimensional codings (or representations) of information and is composed of an encoder \mathcal{E} and decoder \mathcal{D} such that $\mathcal{D}(\mathcal{E}(\mathbf{x})) \approx \mathbf{x}$ with $\mathcal{E}(\mathbf{x})$ being of much lower dimension than \mathbf{x} . The idea is that one is preserving the “essence” of \mathbf{x} in creating $\mathcal{E}(\mathbf{x})$, which can be considered a model for \mathbf{x} and that, ideally, $\mathcal{D}(\mathcal{E}(\mathbf{x}))$ is removing only useless artifacts or noise. A sparsity-regularized autoencoder can be obtained from a set of training data $\{\mathbf{x}_i\}$ as the solution of the optimization

$$\mathcal{E}, \mathcal{D} = \underset{\mathcal{E}, \mathcal{D}}{\operatorname{argmin}} \sum_k \|\mathbf{x}_k - \mathcal{D}(\mathcal{E}(\mathbf{x}_k))\|_2^2 \quad \text{s.t.} \|\mathcal{E}(\mathbf{x}_k)\|_0 \leq K \quad (17)$$

where \mathcal{E} and \mathcal{D} are both NNs whose parameters are learned via solving (17). Once \mathcal{E} and \mathcal{D} of the autoencoder are found, they can be used as a prior to reconstruct an image using an MBIR formulation

$$\begin{aligned} \hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} & \|\mathbf{y} - \mathcal{C}(\mathbf{x})\|_2^2 + \lambda \|\mathbf{x} - \mathcal{D}(\mathcal{E}(\mathbf{x}))\|_2^2 \\ \text{s.t.} & \|\mathcal{E}(\mathbf{x})\|_0 \leq K. \end{aligned} \quad (18)$$

Such an approach was adopted, for example, in [33], which applied (18) using a formulation based on image patches, solving the resulting optimization via an alternating minimization.

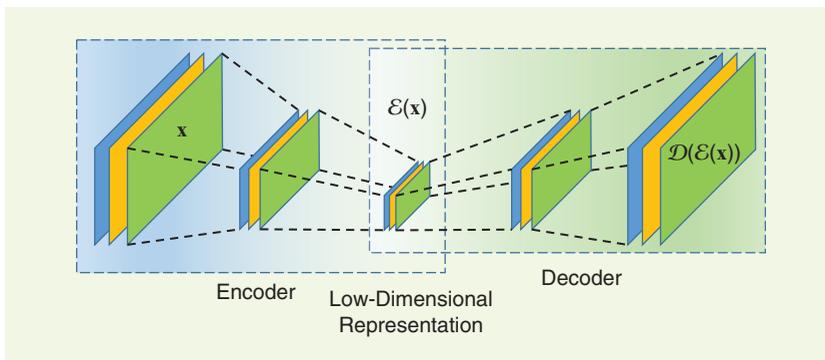


FIGURE 6. A generic autoencoder structure.

Learned inverse mappings

Deep learning can also be used to train a network that directly implements the inverse mapping from the data \mathbf{y} to estimate $\hat{\mathbf{x}}$ (e.g., [34]). That is, we learn a function \mathcal{F} such that $\hat{\mathbf{x}} = \mathcal{F}(\mathbf{y})$. This approach was taken in [35] with a method termed *AUTOMAP*. Four cases motivated by tomography and MRI were considered: tomographic inversion, spiral k -space MRI data, undersampled k -space MRI data, and misaligned k -space MRI data. The AUTOMAP framework used a general feed-forward deep NN architecture composed of fully connected layers followed by a sparse convolutional autoencoder. The network for each application was learned from a set of corresponding training data without the inclusion of any physical or expert knowledge. Taking k -space MRI as an example, although we know that the inverse Fourier transform produces the desired image, the AUTOMAP network has to discover this fact directly from the training data alone. While the results are intriguing, the approach suffers from the large number of training parameters required by the multiple fully connected layers, which seems to limit its application to relatively small problems.

Deep image priors

The final approach that we consider here is the deep image prior (DIP) [36], which replaces the explicit regularizer $\mathcal{R}(\mathbf{x})$ in the MBIR optimization (7) with an implicit prior in the form of an NN. Specifically, as in the autoencoders discussed previously, the DIP NN is a generator or decoder, $\mathcal{D}(\cdot)$, that maps a reduced-dimensional “code” vector \mathbf{w} to a reconstructed image, i.e., $\hat{\mathbf{x}} = \mathcal{D}(\mathbf{w})$. The DIP formulation solves for the decoder so as to minimize the loss with respect to the observation \mathbf{y} , i.e.,

$$\mathcal{D} = \underset{\mathcal{D}}{\operatorname{argmin}} \mathcal{L}(\mathbf{y}, \mathcal{D}(\mathbf{w})). \quad (19)$$

For example, one might use a loss in the form of $\mathcal{L}(\mathbf{y}, \mathcal{D}(\mathbf{w})) = \|\mathbf{y} - \mathcal{C}(\mathcal{D}(\mathbf{w}))\|_2^2$, as in (8). While optimization might also include \mathbf{w} , usually the code \mathbf{w} is chosen at random and kept fixed; additionally, the initial network \mathcal{D} is also typically chosen at random. In essence, (19) imposes an implicit regularization such that $\mathcal{R}(\mathbf{x}) = 0$ for images \mathbf{x} that can be produced by the deep NN \mathcal{D} , and $\mathcal{R}(\mathbf{x}) = \infty$ otherwise. DIP formulations have been applied to a number of image-reconstruction tasks,

including denoising, superresolution, and inpainting [36], as well as in applications such as MRI, computed tomography, and radar imaging, along with other computational-imaging modalities illustrated in Figure 1. The DIP approach has the advantage that the decoder \mathcal{D} is learned during the DIP optimization (19) as applied to the specific observed \mathbf{y} in question; that is, no large body of training data is required. On the other hand, the number of iterations conducted while solving (19) must be carefully controlled so as to avoid the overfitting of $\hat{\mathbf{x}} = \mathcal{D}(\mathbf{w})$ to \mathbf{y} .

Folding learning into algorithms: Coming full circle

In the “Physics-Driven Imaging: Explicit Inversion” section, we noted that the early stages of imaging were dominated by the drive for physics-derived algorithms that achieve the inverse of the physical sensing operator, which is then applied to the observed data. This emphasis on algorithmic inversion was subsequently replaced with the optimization framework of MBIR, allowing more complex sensing configurations and the inclusion of prior information. In the service of increasing the role of learning in computational imaging, researchers have recently shifted their focus from optimization back to algorithms, but this time using learned—rather than analytically derived—elements. We now discuss some examples of recent developments in this vein with an aim to present representative examples of these methods rather than an exhaustive catalog.

Perhaps the most straightforward example of this paradigm is to start with a physically derived inverse algorithm and then replace parts of the original algorithm with alternative learned elements. For example, in [37], [38], the physics-appropriate FBP algorithm of tomography is mapped into an NN architecture wherein the back-projection stage is held fixed while the filter (i.e., the combining weights) is instead learned from data to minimize a loss function.

A particularly impactful and popular approach has been plug-and-play priors (PPPs) [39], [40], which was motivated by the advances in image denoising within the image processing community. Effectively, PPP incorporated into the MBIR image-formation framework the power of advanced denoisers—such as those based on nonlocal means [41] or block-matching and 3D filtering (BM3D) [42]—even though these denoisers are not simple solutions of an underlying optimization problem. Specifically, the PPP framework originated from the ADMM approach [(12), (13), and (14)] for solving the MBIR problem (7). ADMM entails iteration over two main subproblems, with one subproblem (13) involving the observation and sensor loss function $\mathcal{L}(\mathbf{y}, \mathcal{C}(\mathbf{x}))$, while the other (12) involves a proximal operator of the prior (or regularization) term $\mathcal{R}(\mathbf{x})$, which has the form of a MAP denoiser. This insight led to the replacement of the prior proximal operator derived from $\mathcal{R}(\mathbf{x})$ in (12) with alternative state-of-the-art denoisers even though these denoisers might have no corresponding explicit function $\mathcal{R}(\mathbf{x})$. The result was an extremely flexible framework leading to state-of-the-art performance in a very diverse range of applications.

The PPP framework allows one to replace explicit image priors specified by the function $\mathcal{R}(\mathbf{x})$ with powerful image-enhancement mappings, including those potentially learned from training data (e.g., [43]). Notable developments within this framework or inspired by it include general proximal algorithms, regularization by denoising (RED), and multiagent consensus equilibrium (MACE). General proximal algorithms arose when it was demonstrated that while the original development of PPP focused on the ADMM algorithm, the same

approach can be applied to any proximal algorithm [24], including, in particular, the accelerated proximal gradient method or the fast iterative shrinkage/thresholding algorithm (ISTA) [44], which are far more suitable for problems with a nonlinear forward operator and which have allowed the development of online variants of PPP [45], i.e., those that use only a subset of the observations \mathbf{y} , rather than the full set, at each iteration. Additionally, RED was proposed as an alternative formulation for exploiting a denoiser in a way that does have an explicit cost function [46]. While this formulation is tenable in only some special cases [47], it

has proven to be a popular alternative to PPP, with a growing range of applications and extensions.

Finally, MACE [48] extended the original PPP framework in two distinct ways: extending the number of terms (or “agents”) that could be addressed via a formulation that is closely related to ADMM consensus [23, Ch. 7] and introducing a more

theoretically sound interpretation of the approach based on fixed-point algorithms rather than as optimization with unknown or nonexistent regularizers. The general nature of MACE has also been used to include learned models into both the observation side of inverse problems as well as the prior-image information [49], allowing balanced inclusion of both types of constraints—Figure 7 shows an example.

Additionally, interest has grown in a set of techniques that can be collected under the labels “algorithm unrolling,” “deep unrolling,” or “unfolding” [50], as depicted in Figure 8. The idea of these methods is to view a fixed number of iterations of an algorithm as a set of layers (or elements) in a network. The resulting network then performs the steps of the algorithm, while the parameters or steps of the original network are then collectively learned from training data or replaced by learned alternatives. One stated benefit of such an approach is that the resulting overall networks, obtained from underlying optimization algorithms, can be interpreted in ways that typical black-box deep networks with many parameters cannot.

Though the number of works in this spirit is now so large that it is impossible to touch on them all, we briefly mention a few final examples. An early instance was the work in [51], which was based on the ISTA for sparse coding and may be the first to refer to algorithm “unrolling.” In [51], a small fixed number of ISTA iterations are used for the network structure, and the associated linear weights are learned. In another example, [52] starts with a projected gradient-descent method to minimize a regularized optimization problem similar to (8), and then the projector onto the space of regularization constraints in the algorithm is replaced by a learned CNN. In a similar vein, the authors of [28] instead start with a solution to (7) provided by 10 iterations of the Chambolle-Pock primal-dual algorithm, and then the functions of the primal and dual steps are replaced by learned CNNs. Another example is found in [53], which unrolls a half-quadratic splitting

Additionally, the use of training data coupled with high-dimensional models can compensate for imperfect analytical knowledge in a problem.

algorithm for the solution of a blind-deconvolution problem to obtain an associated network with the underlying parameters then being learned. We note that while many more examples are discussed in [50], they all largely retain physics-derived observation models while incorporating prior information implicitly through learning or training processes for algorithmic parameters.

While a variety of NN models have been used in this context, there has been significant recent interest in generative models, including various forms of generative adversarial networks [54], normalizing flows [55], and score-based diffusion models [56], [57]. The stochastic nature of these methods also makes them attractive for uncertainty characterization for computational imaging, another alternative for which is to use Bayesian NNs [58]. Finally, the challenges associated with access to the ground truth required for supervised learning have led to the development of self-supervised or unsupervised methods for learning-based computational imaging [59], [60], an important current direction.

TCI focuses on solutions to imaging problems in which computation plays an integral role in the formation of an image from sensor data.

The rise of learning-based methods discussed in this section represents a great increase in the richness of the prior information that can be included in computational-imaging problems. Additionally, the use of training data coupled with high-dimensional models can compensate for imperfect analytical knowledge in a problem. Consequently, these methods are an ongoing area of active research and promise to impact a wide range of application areas. In our taxonomy of computational-imaging approaches, these techniques are characterized by relatively high computation (for parameter learning) coupled with the use of big data.

Consequently, these methods are an ongoing area of active research and promise to impact a wide range of application areas. In our taxonomy of computational-imaging approaches, these techniques are characterized by relatively high computation (for parameter learning) coupled with the use of big data.

SPS involvement

This section provides a brief history of SPS initiatives and activities related to computational imaging, including the establishment of *IEEE Transactions on Computational Imaging (TCI)* and the SPS Computational Imaging Technical Committee (CI TC) as well as support for community-wide conference and seminar activities.

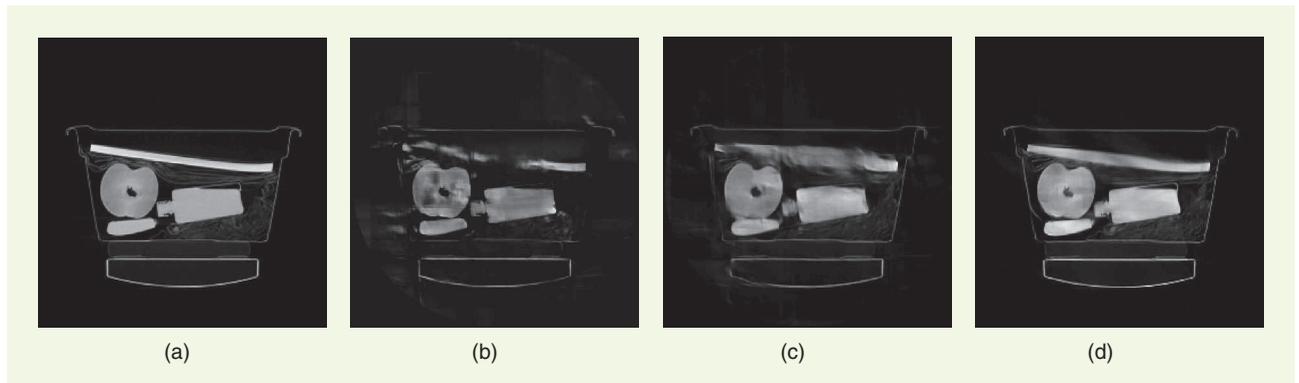


FIGURE 7. Tomographic reconstructions with challenging limited-angle scanning in a baggage-screening security application. (a) Ground truth (FBP reconstruction using complete scanning). (b) FBP reconstruction followed by CNN post-processing using half-scan data. (c) PPP reconstruction with an image-domain learned denoiser using half-scan data. (d) MACE reconstruction including both learned data- and image-domain models using half-scan data. (Source: Adapted from [49], ©2021 IEEE.)

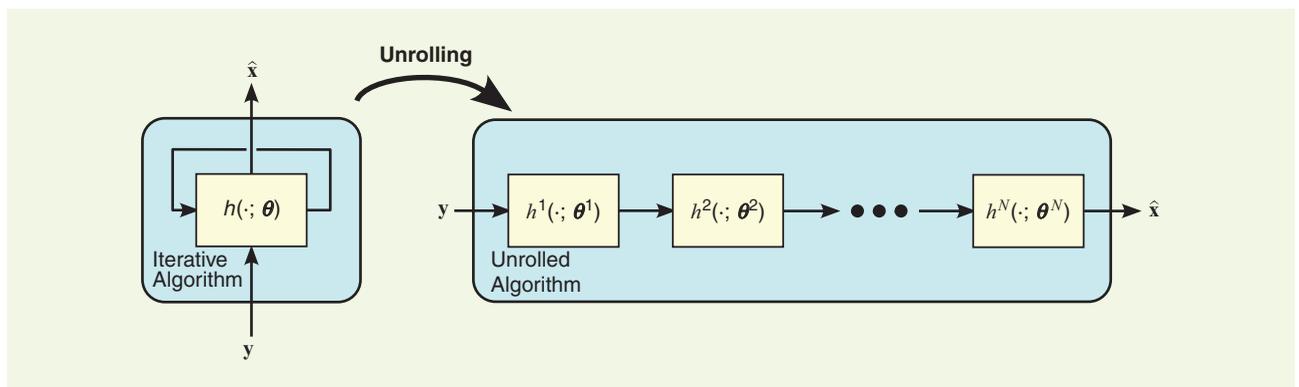


FIGURE 8. A general framework for algorithm unrolling. Left: An iterative algorithm composed of iterations of the fixed function $h(\cdot)$ with parameter set Θ . Right: An unrolling of N iterations into a network with multiple sublayers of h^k , where each h^k now represents a (possibly structured) network with (possibly) different parameters Θ^k ; these networks are then learned through a training process. (Source: Adapted from [50].)

TCI

Motivated by the rapid growth of computational imaging as a research and technology area distinct from image processing, the creation of a new journal on computational imaging was first proposed to the SPS Technical Activities Board Periodicals Committee in 2013 in an effort led by three serving and prior editors-in-chief of *IEEE Transactions on Image Processing*: Charles Bouman, Thrassos Pappas, and Clem Karl. The motivation for the new journal was the rapid growth of computational imaging as a research and technology area distinct from image processing. The journal was launched in 2015, with Clem Karl as its inaugural editor-in-chief.

TCI focuses on solutions to imaging problems in which computation plays an integral role in the formation of an image from sensor data. The journal's scope includes all areas of computational imaging, ranging from theoretical foundations and methods to innovative computational-imaging system design. Topics of interest include advanced algorithms and mathematical methodology, model-based data inversion, methods for image recovery from sparse and incomplete data, techniques for nontraditional sensing of image data, methods for dynamic information acquisition and extraction from imaging sensors, and software and hardware for efficient computation in imaging systems.

TCI has grown rapidly since its inception. With around 40 submissions a month and an impact factor of 4.7, it is now one of the leading venues for the publication of computational-imaging research. *TCI* is somewhat unique within the SPS publications portfolio in that it draws submissions from a broad range of professional communities beyond the SPS, including SIAM, Optica (formerly the Optical Society of America), and SPIE, and that it has connections to a broad range of domains, including radar sensing, X-ray imaging, optical microscopy, and ultrasound sensing. The editorial board similarly includes members from these diverse communities.

CI TC

The Computational Imaging Special Interest Group was established within the SPS in 2015 and was promoted to TC status in 2018. The goal of the CI TC is the promotion of computational imaging as well as the formation of a community of computational-imaging researchers that crosses the traditional boundaries of professional-society affiliations and academic disciplines, with activities including the sponsorship of workshops and special sessions, assistance with the review of papers submitted to the major society conferences, and the support and promotion of *TCI*. The CI TC currently consists of the chair, the vice chair (or the past chair), and 40 regular voting members. Additionally, there are nonvoting advisory members, associate members, and affiliate members. To promote collaboration across communities, the CI TC also appoints liaisons to other IEEE and non-IEEE professional groups whose interests intersect with computational imaging.

The Computational Imaging Special Interest Group was established within the SPS in 2015 and was promoted to TC status in 2018.

Those interested in becoming involved with the CI TC can become an affiliate member via an easy web-based registration process (see <https://signalprocessingsociety.org/community-involvement/computational-imaging/affiliates>). Affiliates are nonelected nonvoting members of the TC, and affiliate membership is open to IEEE Members of all grades as well as to members of certain other professional organizations in interdisciplinary fields within the CI TC's scope.

Vision and outlook

Although computational imaging has only recently emerged as a distinct area of research, it rests upon several decades of work performed in separate research and technology communities. Accordingly, a broad collaboration of researchers—hailing from

signal processing, machine learning, statistics, optimization, optics, and computer vision with domain expertise in various application areas ranging all the way from medical imaging and computational photography to remote sensing—is essential for accelerating progress in this area. This progress is exactly what *TCI* and the CI TC aim to catalyze. Most recently, the community has seen an increasing involvement and impact of machine learning in computational imaging, a theme to which TC members and *TCI* authors make significant contributions.

Conclusions

In conclusion, computational imaging has emerged as a significant and distinct area of research. With the increasing demand for improved information extraction from existing imaging systems coupled with the growth of novel imaging applications and sensing configurations, researchers have turned to advanced computational and algorithmic techniques. These techniques, including deep learning methods, have led to the development of new imaging modalities as well as the ability to process and analyze large datasets. The SPS has played an important role in this growing area through the creation of a new highly ranked journal, a new energetic TC, and support for new cross-society conference and seminar activities. The continued advancement of computational imaging will impact a wide range of applications, from health care to science to defense and beyond.

Authors

W. Clem Karl (wckarl@bu.edu) received his Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology. He is currently the chair of the Electrical and Computer Engineering Department and a member of the Biomedical Engineering Department, Boston University, Boston, MA 02215 USA. He was the inaugural editor-in-chief of *IEEE Transactions on Computational Imaging* and the editor-in-chief of *IEEE Transactions on Image Processing*. He has served in many society roles including on the IEEE Signal Processing Society Board of Governors and the IEEE Signal Processing Society Publications Board. He is a Fellow of IEEE and the American Institute for Medical and Biological Engineering.

James E. Fowler (fowler@ece.msstate.edu) received his Ph.D. degree in electrical engineering from The Ohio State University. He is currently a William L. Giles Distinguished Professor in the Department of Electrical and Computer Engineering, Mississippi State University, Mississippi State, MS 39762 USA. He was previously the editor-in-chief of *IEEE Signal Processing Letters*, a senior area editor for *IEEE Transactions on Image Processing*, and an associate editor for *IEEE Transactions on Computational Imaging*. He is currently the chair of the Computational Imaging Technical Committee of the IEEE Signal Processing Society. He is a Fellow of IEEE.

Charles A. Bouman (bouman@purdue.edu) received his Ph.D. degree from Princeton University. He is the Showalter Professor of Electrical and Computer Engineering and Biomedical Engineering at Purdue University, West Lafayette, IN 47907-1285 USA. He has served as the IEEE Signal Processing Society's vice president of technical directions as well as the editor-in-chief of *IEEE Transactions on Image Processing*. He is a Fellow of IEEE, the American Institute for Medical and Biological Engineering, the Society for Imaging Science and Technology (IS&T), and SPIE and a member of the National Academy of Inventors.

Müjdat Çetin (mujdat.cetin@rochester.edu) received his Ph.D. degree from Boston University. He is a professor of electrical and computer engineering and of computer science, the director of the Goergen Institute for Data Science, and the director of the New York State Center of Excellence in Data Science, University of Rochester, Rochester, NY 14627-0001 USA. He is currently serving as the editor-in-chief of *IEEE Transactions on Computational Imaging* and a senior area editor for *IEEE Transactions on Image Processing*. He previously served as the chair of the IEEE Computational Imaging Technical Committee. He is a Fellow of IEEE.

Brendt Wohlberg (brendt@ieee.org) received his Ph.D. degree in electrical engineering from the University of Cape Town. He is currently a staff scientist with the Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545-1663 USA. He was the chair of the Computational Imaging Special Interest Group (now the Computational Imaging Technical Committee) of the IEEE Signal Processing Society and was the editor-in-chief of *IEEE Transactions on Computational Imaging*. He is currently the editor-in-chief of *IEEE Open Journal of Signal Processing*. He is a Fellow of IEEE.

Jong Chul Ye (jong.ye@kaist.ac.kr) received his Ph.D. degree from Purdue University. He is a professor of the Kim Jaechul Graduate School of Artificial Intelligence and an adjunct professor at the Department of Mathematical Sciences and the Department of Bio and Brain Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon 305-701, Korea. He has served as an associate editor of both *IEEE Transactions on Image Processing* and *IEEE Transactions on Computational Imaging*. He was the chair of the Computational Imaging Technical Committee of the IEEE Signal Processing Society. He is a Fellow of IEEE.

References

- [1] Z. Wang and G. AlRegib, "Interactive fault extraction in 3-D seismic data using the Hough transform and tracking vectors," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 99–109, Mar. 2017, doi: 10.1109/TCL.2016.2626998.
- [2] L. Tian, X. Li, K. Ramchandran, and L. Waller, "Multiplexed coded illumination for Fourier Ptychography with an LED array microscope," *Biomed. Opt. Exp.*, vol. 5, no. 7, pp. 2376–2389, Jul. 2014, doi: 10.1364/BOE.5.002376.
- [3] K. L. Bouman, "Portrait of a black hole: Here's how the event horizon telescope team pieced together a now-famous image," *IEEE Spectr.*, vol. 57, no. 2, pp. 22–29, Feb. 2020, doi: 10.1109/MSPEC.2020.8976898.
- [4] L. A. Shepp and B. F. Logan, "The Fourier reconstruction of a head section," *IEEE Trans. Nucl. Sci.*, vol. 21, no. 3, pp. 21–43, Jun. 1974, doi: 10.1109/TNS.1974.6499235. [Online]. Available: <http://ieeexplore.ieee.org/document/6499235/>
- [5] Z.-P. Liang and P. C. Lauterbur, *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. New York, NY, USA: IEEE Press, 2000.
- [6] C. V. Jakowatz, D. E. Wahl, P. A. Thompson, P. H. Eichel, and D. C. Ghiglia, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. New York, NY, USA: Springer, 1996.
- [7] M. Vauhkonen, D. Vadasz, P. Karjalainen, E. Somersalo, and J. Kaipio, "Tikhonov regularization and prior information in electrical impedance tomography," *IEEE Trans. Med. Imag.*, vol. 17, no. 2, pp. 285–293, Apr. 1998, doi: 10.1109/42.700740. [Online]. Available: <http://ieeexplore.ieee.org/document/700740/>
- [8] J. Besag, "Digital image processing: Towards Bayesian image analysis," *J. Appl. Statist.*, vol. 16, no. 3, pp. 395–407, Jan. 1989, doi: 10.1080/02664768900000049.
- [9] S. Geman and D. E. McClure, "Statistical methods for tomographic image reconstruction," in *Proc. 46th Session Int. Statist. Inst.*, 1987, vol. 52, pp. 5–21.
- [10] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Process.*, vol. 6, no. 2, pp. 298–311, Feb. 1997, doi: 10.1109/83.551699. [Online]. Available: <https://ieeexplore.ieee.org/document/551699/>
- [11] P. Green, "Bayesian reconstructions from emission tomography data using a modified EM algorithm," *IEEE Trans. Med. Imag.*, vol. 9, no. 1, pp. 84–93, Mar. 1990, doi: 10.1109/42.52985. [Online]. Available: <http://ieeexplore.ieee.org/document/52985/>
- [12] A. Blake, "Comparison of the efficiency of deterministic and stochastic algorithms for visual reconstruction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 1, pp. 2–12, Jan. 1989, doi: 10.1109/34.23109. [Online]. Available: <http://ieeexplore.ieee.org/document/23109/>
- [13] P. J. Huber, "Robust estimation of a location parameter," *Ann. Math. Statist.*, vol. 35, no. 1, pp. 73–101, Mar. 1964, doi: 10.1214/aoms/1177703732. [Online]. Available: <http://projecteuclid.org/euclid.aoms/1177703732>
- [14] T. Hebert and R. Leahy, "A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors," *IEEE Trans. Med. Imag.*, vol. 8, no. 2, pp. 194–202, Jun. 1989, doi: 10.1109/42.24868. [Online]. Available: <http://ieeexplore.ieee.org/document/24868/>
- [15] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 3, pp. 367–383, Mar. 1992, doi: 10.1109/34.120331. [Online]. Available: <http://ieeexplore.ieee.org/document/120331/>
- [16] C. Bouman and K. Sauer, "A generalized Gaussian image model for edge-preserving MAP estimation," *IEEE Trans. Image Process.*, vol. 2, no. 3, pp. 296–310, Jul. 1993, doi: 10.1109/83.236536. [Online]. Available: <http://ieeexplore.ieee.org/document/236536/>
- [17] J. Mairal, "Sparse modeling for image and vision processing," *Found. Trends® Comput. Graphics Vision*, vol. 8, nos. 2–3, pp. 85–283, Dec. 2014, doi: 10.1561/06000000058. [Online]. Available: <http://www.nowpublishers.com/articles/foundations-and-trends-in-computer-graphics-and-vision/CGV-058>
- [18] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006, doi: 10.1109/TIT.2006.871582. [Online]. Available: <http://ieeexplore.ieee.org/document/1614066/>
- [19] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006, doi: 10.1109/TIT.2005.862083. [Online]. Available: <http://ieeexplore.ieee.org/document/1580791/>
- [20] D. Malioutov, M. Cetin, and A. Willsky, "Optimal sparse representations in general overcomplete bases," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Montreal, QC, Canada: IEEE, 2004, vol. 2, pp. ii–793–6, doi: 10.1109/ICASSP.2004.1326377.
- [21] S. Mun and J. E. Fowler, "Block compressed sensing of images using directional transforms," in *Proc. IEEE Int. Conf. Image Process.*, Cairo, Egypt, Nov. 2009, pp. 3021–3024, doi: 10.1109/ICIP.2009.5414429.
- [22] D. Takhar, J. N. Laska, M. B. Wakin, M. F. Duarte, D. Baron, S. Sarvotham, K. F. Kelly, and R. G. Baraniuk, "A new compressive imaging camera architecture

- using optical-domain compression,” in *Proc. Comput. Imag. IV (SPIE)*, San Jose, CA, USA, Jan. 2006, p. 606509, doi: 10.1117/12.659602.
- [23] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jul. 2011, doi: 10.1561/2200000016. [Online]. Available: <https://www.nowpublishers.com/article/Details/MAL-016>
- [24] N. Parikh and S. Boyd, “Proximal algorithms,” *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, Jan. 2014, doi: 10.1561/2400000003.
- [25] P. L. Combettes and J.-C. Pesquet, “Proximal splitting methods in signal processing,” in *Proc. Fixed-Point Algorithms Inverse Probl. Sci. Eng.*, H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, Eds. New York, NY, USA: Springer, 2011, vol. 49, pp. 185–212.
- [26] B. A. Olshausen and D. J. Field, “Emergence of simple-cell receptive field properties by learning a sparse code for natural images,” *Nature*, vol. 381, no. 6583, pp. 607–609, Jun. 1996, doi: 10.11038/381607a0. [Online]. Available: <http://www.nature.com/articles/381607a0>
- [27] M. Elad and M. Aharon, “Image denoising via sparse and redundant representations over learned dictionaries,” *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, Dec. 2006, doi: 10.1109/TIP.2006.881969. [Online]. Available: <http://ieeexplore.ieee.org/document/4011956/>
- [28] J. Adler and O. Oktem, “Learned primal-dual reconstruction,” *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1322–1332, Jun. 2018, doi: 10.1109/TMI.2018.2799231. [Online]. Available: <https://ieeexplore.ieee.org/document/8271999/>
- [29] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, “Deep convolutional neural network for inverse problems in imaging,” *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, Sep. 2017, doi: 10.1109/TIP.2017.2713099. [Online]. Available: <http://ieeexplore.ieee.org/document/7949028/>
- [30] C. You et al., “Structurally-sensitive multi-scale deep neural network for low-dose CT denoising,” *IEEE Access*, vol. 6, pp. 41,839–41,855, Jul. 2018, doi: 10.1109/ACCESS.2018.2858196.
- [31] G. Yang et al., “DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction,” *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1310–1321, Jun. 2018, doi: 10.1109/TMI.2017.2785879. [Online]. Available: <https://ieeexplore.ieee.org/document/8233175/>
- [32] M. U. Ghani and W. C. Karl, “Deep learning-based sinogram completion for low-dose CT,” in *Proc. IEEE 13th Image, Video, Multidimensional Signal Process. Workshop (IVMSP)*, Zagorochoria, Greece: IEEE, Jun. 2018, pp. 1–5, doi: 10.1109/IVMSPW.2018.8448403.
- [33] D. Wu, K. Kim, G. El Fakhri, and Q. Li, “Iterative low-dose CT reconstruction with priors trained by artificial neural network,” *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2479–2486, Dec. 2017, doi: 10.1109/TMI.2017.2753138. [Online]. Available: <https://ieeexplore.ieee.org/document/8038851/>
- [34] Y. Wu and Y. Lin, “InversionNet: An efficient and accurate data-driven full waveform inversion,” *IEEE Trans. Comput. Imag.*, vol. 6, pp. 419–433, 2020, doi: 10.1109/TCI.2019.2956866. [Online]. Available: <https://ieeexplore.ieee.org/document/8918045>
- [35] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, “Image reconstruction by domain-transform manifold learning,” *Nature*, vol. 555, no. 7697, pp. 487–492, Mar. 2018, doi: 10.1038/nature25988. [Online]. Available: <http://www.nature.com/articles/nature25988>
- [36] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior,” *Int. J. Comput. Vision*, vol. 128, pp. 1867–1888, Jul. 2020, doi: 10.1007/s11263-020-01303-4.
- [37] T. Würfl, F. C. Ghesu, V. Christlein, and A. Maier, “Deep learning computed tomography,” in *Proc. Med. Image Comput. Comput.-Assisted Intervention (MICCAI)*, S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, and W. Wells, Eds. Cham, Switzerland: Springer International Publishing, 2016, vol. 9902, pp. 432–440.
- [38] D. H. Ye, G. T. Buzzard, M. Ruby, and C. A. Bouman, “Deep back projection for sparse-view CT reconstruction,” in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Anaheim, CA, USA: IEEE, Nov. 2018, pp. 1–5, doi: 10.1109/GlobalSIP.2018.8646669.
- [39] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, “Plug-and-play priors for model based reconstruction,” in *Proc. IEEE Global Conf. Signal Inf. Process.*, Dec. 2013, pp. 945–948, doi: 10.1109/GlobalSIP.2013.6737048.
- [40] S. Sreehari, S. V. Venkatakrishnan, B. Wohlberg, G. T. Buzzard, L. F. Drummy, J. P. Simmons, and C. A. Bouman, “Plug-and-play priors for bright field electron tomography and sparse interpolation,” *IEEE Trans. Comput. Imag.*, vol. 2, no. 4, pp. 408–423, Dec. 2016, doi: 10.1109/TCI.2016.2599778. [Online]. Available: <http://ieeexplore.ieee.org/document/7542195/>
- [41] A. Buades, B. Coll, and J.-M. Morel, “Non-local means denoising,” *IPOL J. Image Process. On Line*, vol. 1, pp. 208–212, Sep. 2011, doi: 10.5201/ipl.2011.bcm_nlm. [Online]. Available: https://www.ipol.im/pub/art/2011/bcm_nlm/?utm_source=doi
- [42] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image restoration by sparse 3D transform-domain collaborative filtering,” in *Proc. Image Process., Algorithms Syst. VI*, J. T. Astola, K. O. Egiazarian, and E. R. Dougherty, Eds. San Jose, CA, USA, Feb. 2008, p. 681207.
- [43] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, “Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising,” *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017, doi: 10.1109/TSP.2017.2662206. [Online]. Available: <https://ieeexplore.ieee.org/document/7839189/>
- [44] U. S. Kamilov, H. Mansour, and B. Wohlberg, “A plug-and-play priors approach for solving nonlinear imaging inverse problems,” *IEEE Signal Process. Lett.*, vol. 24, no. 12, pp. 1872–1876, Dec. 2017, doi: 10.1109/LSP.2017.2763583. [Online]. Available: <http://ieeexplore.ieee.org/document/8068267/>
- [45] Y. Sun, B. Wohlberg, and U. S. Kamilov, “An online plug-and-play algorithm for regularized image reconstruction,” *IEEE Trans. Comput. Imag.*, vol. 5, no. 3, pp. 395–408, Sep. 2019, doi: 10.1109/TCI.2019.2893568.
- [46] Y. Romano, M. Elad, and P. Milanfar, “The little engine that could: Regularization by denoising (RED),” *SIAM J. Imag. Sci.*, vol. 10, no. 4, pp. 1804–1844, Jan. 2017, doi: 10.1137/16M1102884.
- [47] E. T. Reehorst and P. Schniter, “Regularization by denoising: Clarifications and new interpretations,” *IEEE Trans. Comput. Imag.*, vol. 5, no. 1, pp. 52–67, Mar. 2019, doi: 10.1109/TCI.2018.2880326. [Online]. Available: <https://ieeexplore.ieee.org/document/8528509/>
- [48] G. T. Buzzard, S. H. Chan, S. Sreehari, and C. A. Bouman, “Plug-and-play unlogged: Optimization-free reconstruction using consensus equilibrium,” *SIAM J. Imag. Sci.*, vol. 11, no. 3, pp. 2001–2020, Jan. 2018, doi: 10.1137/17M1122451.
- [49] M. U. Ghani and W. C. Karl, “Data and image prior integration for image reconstruction using consensus equilibrium,” *IEEE Trans. Comput. Imag.*, vol. 7, pp. 297–308, Mar. 2021, doi: 10.1109/TCI.2021.3062986. [Online]. Available: <https://ieeexplore.ieee.org/document/9366922/>
- [50] V. Monga, Y. Li, and Y. C. Eldar, “Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing,” *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 18–44, Mar. 2021, doi: 10.1109/MSP.2020.3016905. [Online]. Available: <https://ieeexplore.ieee.org/document/9363511/>
- [51] K. Gregor and Y. LeCun, “Learning fast approximations of sparse coding,” in *Proc. 27th Int. Conf. Mach. Learn.*, Haifa, Israel: Omnipress, Jun. 2010, pp. 399–406.
- [52] H. Gupta, K. H. Jin, H. Q. Nguyen, M. T. McCann, and M. Unser, “CNN-based projected gradient descent for consistent CT image reconstruction,” *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1440–1453, Jun. 2018, doi: 10.1109/TMI.2018.2832656. [Online]. Available: <https://ieeexplore.ieee.org/document/8353870/>
- [53] Y. Li, M. Tofighi, V. Monga, and Y. C. Eldar, “An algorithm unrolling approach to deep image deblurring,” in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Brighton, U.K.: IEEE, May 2019, pp. 7675–7679, doi: 10.1109/ICASSP.2019.8682542.
- [54] S. Lunz, O. Öktem, and C.-B. Schönlieb, “Adversarial regularizers in inverse problems,” in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, Red Hook, NY, USA: Curran Associates, Inc., Dec. 2018, pp. 8516–8525.
- [55] H. Sun and K. L. Bouman, “Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging,” in *Proc. AAAI Conf. Artif. Intell.*, May 2021, vol. 35, no. 3, pp. 2628–2637, doi: 10.1609/aaai.v35i3.16366. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/16366>
- [56] Y. Song, L. Shen, L. Xing, and S. Ermon, “Solving inverse problems in medical imaging with score-based generative models,” in *Proc. Int. Conf. Learn. Representations*, Mar. 2022. [Online]. Available: <https://openreview.net/forum?id=vaRCHVj0uGI>
- [57] H. Chung and J. C. Ye, “Score-based diffusion models for accelerated MRI,” *Med. Image Anal.*, vol. 80, Aug. 2022, Art. no. 102479, doi: 10.1016/j.media.2022.102479. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1361841522001268>
- [58] C. Ekmekci and M. Cetin, “Uncertainty quantification for deep unrolling-based computational imaging,” *IEEE Trans. Comput. Imag.*, vol. 8, pp. 1195–1209, Dec. 2022, doi: 10.1109/TCI.2022.3233185.
- [59] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, “Noise2Noise: Learning image restoration without clean data,” in *Proc. 35th Int. Conf. Mach. Learn.*, PMLR, Jul. 2018, pp. 2965–2974. [Online]. Available: <https://proceedings.mlr.press/v80/lehtinen18a.html>
- [60] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, “RARE: Image reconstruction using deep priors learned without groundtruth,” *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 6, pp. 1088–1099, Oct. 2020, doi: 10.1109/JSTSP.2020.2998402. [Online]. Available: <https://ieeexplore.ieee.org/document/9103213/>